LEVEL II

COMPUTER SCIENCE
TECHNICAL REPORT SERIES

UNIVERSITY OF MARYLAND
COLLEGE PARK, MARYLAND
20742

LEVEL II ①

TR-1009
DAAG-53-76C-0138

February 1981

DETERMINING THE INSTANTANEOUS DIRECTION OF MOTION FROM OPTICAL FLOW GENERATED BY A CURVILINEARLY MOVING OBSERVER

K. Prazdny

Computer Vision Laboratory
Computer Science Center
University of Maryland
College Park, MD 20742

ABSTRACT

A method is described capable of decomposing the optical flow into its rotational and translational components. The translational component is extracted implicitly by locating the focus of expansion associated with the translational component of the relative motion. The method is simple, relying on minimizing an (error) function of 3 parameters. As such, it can also be applied, without modification, in the case of noisy input information. Unlike the previous attempts at interpreting optical flow to obtain information about the three-dimensional disposition of texture elements, the method uses only relationships between quantities on the projection plane. No 3D geometry is involved. Also outlined is a possible use of the method for the extraction of that part of the optical flow containing information about relative depth directly from the image intensity values, without extracting the "retinal" velocity vectors.

The support of the Defense Advanced Research Projects Agency and the U.S. Army Night Vision Laboratory under Contract DAAG-53-76C-0138 (DARPA order 3206) is gratefully acknowledged, as is the help of Janet Salzman in preparing this paper.

The author also thanks Professor Azriel Rosenfeld for his useful comments.

1. Introduction

The distribution of velocities on the projection surface arising as a consequence of the relative motion of objects with respect to the observer, the optical flow (Gibson, 1950, 1955), contains information not only about the (relative) motion itself, but also about the three-dimensional disposition of the set of texture points of which a given set of image elements is a projection <Note 1>. The distribution of image velocities on the projection surface is a function of three parameters: the (relative) motion of objects, their distance to the center of projection, and the local three-dimensional geometry of the objects. Fortunately, however, their effects are, conceptually at least, separable (Prazdny, 1981).

The problem of extracting information from optical flow can conveniently be divided into two stages. First, one has to develop a method of extracting the image velocities from changes in the image intensity values on the projection surfaces. Following this, it remains to solve the problem of computing the required information about (local) surface orientation, relative depth and motion from the distribution of these velocities. While the interpretation of optical flows logically depends on the solution to the problem of extracting the constituent velocities, the problem can and should be studied on its own, for its solution not only explicitly voices the requirements on the quality of the velocity extraction process, but also determines the ultimate success of the whole enterprise.

Recently, optical flows (and time varying imagery in general) have received growing attention among the computer vision community as a source of possible information about a scene. Nakayama and Loomis (1974) and Fennema and Thompson (1979) studied how discontinuities in the "retinal" velocity field could be used for segmentation purposes. Clocksin (1980), Gibson et al. (1955), and Lee (1974) studied how the optical flow generated by an observer translating in a stationary world provides information about (local) surface orientation. Koenderink and van Doorn (1976), Longuet-Higgins and Prazdny (1980), and Prazdny (1980, 1981) studied the extraction of surface orientation, (relative) depth and motion from optical flow generated by an arbitrary curvilinear motion. Another kind of approach, sometimes considered more suitable for a computer vision system, relies on interpreting a sequence of static images as discrete snapshots (see e.g., Nagel, 1981; Aggarwal and Badler, 1980). The computation of "retinal" velocities from image intensity values was studied by Fennema and Thompson (1979), Hadani et al. (1980), and Horn and Schunck (1980). Similar approaches based on matching various higher-order image structures obtained from two (temporally) consecutive images were attempted by, for example, Barnard and Thompson (1980).

In this paper, we outline a method for obtaining the instantaneous direction of (relative) motion from optical flow manifested as image motions on the planar projection surface. We use polar projection as the model of the physical image forming process. Also, we assume throughout the paper that our world contains only rigid and opaque objects. The method presented here does not use projective or geometrical relations, as might be expected from the use of the polar projection; it is based on computations and relationships defined and measurable solely on the projection plane. For example, the method does not require knowing the visual direction (a 3D vector) of a "retinal" point (as was required, e.g., in Prazdny, 1980).

Before outlining the method, we briefly consider a few relevant facts. Optical flow can be (instantaneously) decomposed into two independent components (Koenderink and van Doorn, 1976; Nakayama and Loomis, 1974), a rotational and a translational component. However, only the translational component contains infromation about (local) surface orientation or relative depth. The translational "retinal" field consists of motion along straight lines all intersecting at a common point, the focus of expansion (FOE). This point corresponds to the point where the (three-dimensional) vector specifying the instantaneous direction of motion (the vector tangent to the motion path described by O at a given instant) pierces the projection plane. Our method, by searching for this point, effectively decomposes the optical flow field into its two constituent fields. Briefly, the method is based on minimizing an error function of three parameters. The construction of the function reflects the following observation: if the three parameters specyifying the rotational component of the (relative) motion are chosen properly, the translational "retinal" field yields lines all meeting at FOE. We do not require the spatial derivatives of the "retinal" velocity field, as in Koenderink and van Doorn (1976) or Longuet-Higgins and Prazdny (1980). The method can, but does not have to be, implemented as a local computation. While we have chosen, for simplicity, to consider only the case of an observer moving in a stationary world, it should be noted that the method has a much more far-reaching implication. Because it produces a description of relative motion, it can be applied to a region of the image locally to describe the (relative) motion of that region independently.

2. Locating the focus of expansion (FOE)

To see that it is possible to decompose the instantaneous positional velocity field on the projection plane into its two components, consider the effects of rotation and translation separately. It is advantageous to imagine that the optical flow field is generated by the motion of the observer in a stationary environment. This conceptualization has an immediate interpretation and is, of course, legitimate, for all motion considered here is relative.

Consider the observer rotation first. Because the rotational component of the relative motion does not carry information about the 3D disposition of the texture elements, the motion of an image element on the projection plane will depend only on its position on PP. A rotation vector (angular velocity vector perpendicular to the instantaneous plane of rotation) can be decomposed into two components, one parallel to the projection plane (PP), and one perpendicular to it (see Figure 1) <Note 2>.

Consider the rotation about the vector perpendicular to PP first (the z-axis in Figure 1). For each "retinal" point P with coordinates (x,y), the rotation of the observer about the z-axis (perpendicular to PP) results in P moving along a circular trajectory on PP. The motion of P on PP is specified by a direction vector $\bar{c}=(-y,x)/\sqrt{(x^2+y^2)}$, and by the magnitude $c=c_0\sqrt{(x^2+y^2)}$, where c_0 is the speed of a "retinal" point at a unit distance from O' (the center of the "retinal" coordinate frame). The (2D) velocity of P(x,y) due to observer rotation about the z-axis is thus given by

$$(1) \quad \underline{c}(x,y)=c_0(-y,x)$$

Consider now the situation in which the observer rotates about a vector parallel to the projection plane <Note 3>. To simplify the discussion, we consider only rotation about an axis (through O) parallel to the "retinal" y-axis. The expression for rotation about a parallel to the x-axis is symmetrical in the coordinates x and y [compare equations (5) and (6)]. We first show that the path of a point P(x,y) under rotation of the observer about the y-axis is a hyperbola, and then derive the expression for the velocity vector $\underline{h}_H$ at P(x,y).

Consider Figure 2. A stationary texture element in the 3D environment projects into a point P(x,y) on PP. As PP rotates about a line parallel to the y-axis, the coordinates of P will eventually become $P(0,y_0)$. Observe that the projecting ray defines a fixed visual angle with respect to the plane of rotation. It is clear from Figure 2 that

$$y_0^2=\tan^2\varepsilon=\frac{y^2}{x^2+1}$$

This is because the distance $|OO'|=1$, by assumption (this effectively scales the whole projection system by the focal distance). From this, we obtain

$$(2) \quad \frac{y^2}{y_0^2} - x^2 = 1$$

This is the equation of a hyperbola with center at origin O'. The direction of the velocity vector at P(x,y) is determined by the tangent line at that point. Differentiating (2), we obtain

$$y' = \frac{xy}{x^2+1} = \tan\xi$$

(see Figure 3). $\bar{h}_H$ is thus determined by $\bar{h}_H = (\cos\xi, \sin\xi)$. In terms of the ("retinal") coordinates (x,y) of P, this becomes

$$(3) \qquad \bar{h}_H = \frac{1}{((x^2+1)+(xy)^2)^{\frac{1}{2}}} \quad (x^2+1,\ xy)$$

To determine the magnitude of $\underline{h}_H$, consider two fixed points R and S on two rays such that at time t_0, the points coincide with points $P(0,y_0)$ and O', respectively (see Figure 3). The two rays define a visual angle ξ. Now at a time t_1 (after a rotation of PP by some angle), R projects into the point (x,y) while S projects into the point (x,0). It is evident that the two projections move so that at any time, their x-coordinates are the same. In other words, the x-components of their velocities on PP are the same. We know that the path of P is a hyperbola. It is thus sufficient to compute the horizontal velocity component and project it back onto $\bar{h}_H$ to obtain h_H, the magnitude of $\underline{h}_H$.

Consider Figure 4. If the point x moves with angular velocity h_0 (recall that $|OO'|=1$), then h_x is defined by

$$h_x = \frac{h_0\sqrt{x^2+1}}{\cos\eta}$$

But $\cos\eta = 1/\sqrt{(x^2+1)}$ (see Figure 4) so that $h_x = h_0(x^2+1)$. Projecting h_x back on $\bar{h}_H$ and combining the result with equation (3) we obtain

$$(5) \qquad \underline{h}_H = h_{0H}\ (x^2+1,\ xy)$$

Here h_{0H} is the speed of a "retinal" element at 0'. Analogously, the rotation of the observer about an axis (through O) parallel to the "retinal" x-axis results in the velocity vector h_V defined by

$$(6) \quad \underline{h}_V = h_{0V}\,(xy,\ y^2+1)$$

The input data we are trying to interpret, the optical flows, consist of a set of vectors $\underline{v}''$ defining the positional velocity field on the projection plane. Because we are dealing with velocities, it is easy to see that

$$(7) \quad \underline{v}'' = \underline{c} + (\underline{h}_V + \underline{h}_H) + \underline{t}$$

where $\underline{t}$ is the velocity vector due to the pure translation of the observer. In other words, for each "retinal" locus (x,y), and a set of parameters (h_{0H}, c_0, h_{0V}), equation (7) defines a vector $\underline{t}$ which is a vector function of the three parameters.

As mentioned above, the property of the translational "retinal" velocity field (defined as straight lines specified by a given "retinal" locus (x,y) and the associated vector $\underline{t}$) is that all the lines intersect at one common point, the FOE (see Figure 5). This property makes it possible to define an error function which will lead to resolution of the vector field $\underline{v}''$ into its rotational and translational components. For a given distribution of $\underline{v}''$ on PP, we are searching for those values of the parameters (h_{0H}, c_0, h_{0V}) for which the set $T=\{\underline{t}_i\}$ is such that all lines L_i defined by the vector $\underline{t}_i$ and the retinal locations P_i intersect at a common point, the FOE.

One way of doing this is as follows. Consider an arbitrary "retinal" point P [with coordinates (x,y)] and a set of other (possibly neighboring) points $\{P_i\}$. The points P_i together with the vectors $\underline{t}_i$ define lines L_i which intersect the line L at points I_i (see Figure 6). Consider the lengths ℓ_i between the intersections I_i and the point P. The variance $V=\sum_i(\ell_i-\tilde{\ell})^2/n$ (where $\tilde{\ell}$ is the algebraic average of the ℓ_is) is a good measure of the dispersion of the intersections I_i. When the lines L_i all meet at FOE, V=0. To obtain the FOE, we thus simply minimize $V(h_{0H},c_0,h_{0V})$. Note the way in which the decomposition is accomplished: a property of the translational field is here used to obtain the rotational field, resulting in both fields being obtained at the same time, by the very same computation. The method, being minimalization of a distribution measure, can also immediately be applied when the input data (the vectors $\underline{v}$") are noisy.

3. Some experimental results

The schema described above was tested in a (simulated) world of planar surfaces. The results are encouraging. Eight points surrounding a central point were used to define the set $\{P_i\}$ to obtain the variance V. It should be noted that while in our implementation neighboring points were used (the neighborhood subtended about 15 degrees of arc), this is by no means a necessary condition. A direct minimalization scheme attributed to Nelder and Mead (Nash, 1979) was used to minimize the variance V. The scheme was used mainly for its simplicity and ease of encoding. The values of h_{0H}, c_0, and h_{0V} were restricted to lie between ±90 degrees of arc/sec (the "negative" values corresponding to counterclockwise rotations). This feasible region was defined to restrict the search space to meaningful values and to prevent possible divergence of the iterative process. The minimalization procedure converged to a correct solution from any initial guess within this feasible region. Not all eight distances ℓ_i were used to define V. To minimize the influence of (quantization) errors, the lengths ℓ_i were ordered in magnitude and the two extremal magnitudes were discarded. We also tried to use the *range* of ℓ_i (defined as $|\ell_{max}-\ell_{min}|$) as the error function with good results. In both cases, the FOE was located precisely (using single precision arithmetic [7 significant digits]). When the precision with which the vectors $\underline{v}''$ were defined was lowered to 4 significant digits (the angular error made by this quantization depends on the magnitude of the vector v" [see Figure 7]), the FOE was located within approximately ±5 degrees of arc of the correct position. Extensive testing with real data (and using a more efficient and faster minimalization schema) should be performed to determine how the errors in $\underline{v}''$ propagate through the computations and affect the precision with which the FOE can be obtained.

4. Discussion and conclusions

It is important to realize precisely what has been achieved, and how. Given a set of "retinal" vectors $\underline{v}$" on the planar projection surface, we have shown that it is possible to extract the translational velocity field, containing all information about spatial disposition of the texture elements, solely by computations using data available on the projection surface (see Figure 8). In fact, besides the velocity vectors $\underline{v}$", only the positions of the corresponding loci on the plane with respect to a fixed reference point (the "fovea") are required. Another feature of the method is that it can be implemented as a local computation (the radius of the neighborhood would have to be large enough), and thus performed at many "retinal" locations in parallel, thus decreasing the dependence of the method on a good initial approximation to the parameters h_{0H}, c_0, and h_{0V}. The simplicity of the method is striking, especially in comparison with other methods purporting to achieve the same results (e.g., Prazdny 1980; see also Nagel 1981). The method requires only a few points and the corresponding "retinal" velocities as input (for example, in the visual periphery, which is apparently used by the human visual system to compute egomotion [Johansson, 1977]). One disadvantage of the method is that it fails when the direction of instantaneous motion is parallel to PP. In this case, the FOE is undefined (it corresponds to an ideal point of the projective plane). This is not a serious drawback, however. Another similar method based on maximizing the parallelism between the vectors defining the translational field could take care of this situation.

It is also important to realize that once the FOE has been computed, we immediately know the direction of the translatory motion on the projection plane at each "retinal" locus; it is simply the line connecting the FOE with the given locus (on the "retinal" plane). To obtain information about (relative) depth or (local) surface orientation, we need to compute only the magnitude of motion in this direction; the two-dimensional problem is thus reduced to a more manageable one-dimensional problem. This leads directly to a much more general schema where only the velocities ($\underline{v}$") of a few "interesting" image elements (at "prominent" locations where the velocity $\underline{v}$" can easily be detected) are computed first to locate the FOE. Following this the magnitude of the translatory motion at each image point would be computed without explicit extraction of the optical flow (the velocities $\underline{v}$") itself. As was noted by Batali and Ullman (1979) or Horn and Schunck (1980), one can compute, by a local computation, only the velocity component in the direction of the gradient of the image intensity function at a given "retinal" locus. But this is all we would need if the FOE were already located: by projecting this velocity component onto the direction of the translatory field at a given image plane locus, we would obtain the (relative) depth information in its purest form - as the distribution of the magnitudes of the translatory field.

To summarize, we have shown how the direction of (relative) motion can be computed by a simple minimization computation operating on data available on the projection surface. The method can be implemented locally and is also feasible biologically. Speculatively, perhaps, its operation might be reflected in the recent finding of the "looming" or changing-size channels in the human visual pathway (Regan, Beverley, and Cynader, 1979; Beverley and Regan, 1979).

Notes

<Note 1>

In general, only conclusions about relative quantities can be derived by interpreting optical flows. Local surface orientation, relative depth (the ratio of distances of two texture elements in two distinct visual directions), and relative motion are such quantities.

<Note 2>

The following notational convention will be used throughout the paper. $\underline{n}$ denotes a vector, $\bar{n}$ is its unit vector, and n is the norm of n, i.e., $\underline{n}=n\bar{n}$. Angular velocities are conceptualized as axial vectors, i.e., vectors perpendicular to the instantaneous plane of rotation, with magnitude equivalent to the angular speed. The word "retinal" will be used to denote the projection plane, PP. P(x,y) denotes a "retinal" point P with "retinal" coordinates (x,y) in the two-dimensional coordinate frame centered at O'.

<Note 3>

The set of paths traced by the image elements under this motion is a family of hyperbolas with principal axes inclined at angle ω with respect to the y-axis. The family is symmetrical about a straight line through O'. This is the line corresponding to the intersection of the plane of rotation with the projection plane.

References

Aggarwal, J. K. and Badler, N. I. (eds.) 1980
Special issue on motion and time-varying imagery
IEEE Trans. Pattern Analysis Machine Intelligence 2,6.

Batali, J. and Ullman, S. 1979
Motion detection and analysis
DARPA Image Understanding Workshop, 69-75
Science Applications Inc., Arlington, VA.

Barnard, S. T. and Thompson, W. B. 1980
Disparity analysis of images
IEEE Trans. Pattern Analysis Machine Intelligence 2,4, 333-340.

Beverley, K. and Regan, D. 1979
Separable aftereffects of changing size and motion in depth
Vis. Res. 19, 727-732.

Clocksin, W. F. 1980
Perception of surface slant and edge labels from optical flow
Perception 9, 253-269.

Fennema, C. L. and Thompson, W. B. 1979
Velocity determination in scenes containing several moving objects
Computer Graphics Image Processing 9, 301-315.

Gibson, J. J. Olum, P. and Rosenblatt, F. 1955
Parallax and perspective during aircraft landing
Am. J. Psych. 68, 372-385.

Hadani, I. Ishai, G. and Gur, M. 1980
Visual stability and space perception in monocular vision
J. Opt. Soc. Am. 70, 60-65.

Johansson, G. 1977
Studies in visual perception of locomotion
Perception 6, 365-376.

Koenderink, J. J. and van Doorn, A. J. 1976
Local structure of movement parallax of the plane
J. Opt. Soc. Am. 66, 717-723.

Lee, D. N. 1974
Visual information during locomotion
in: Perception: Essays in Honor of J. J. Gibson
eds. McLeod, R. B. and Pick, H. L.
Ithaca, NY: Cornell University Press.

Longuet-Higgins C. and Prazdny, K. 1980
The interpretation of a moving retinal image
Proc. R. Soc. London B-208, 385-397.

Nagel, H. H. 1981
On the derivation of 3D rigid point configurations from image sequences
IEEE Conf. Pattern Recognition Image Processing
(to appear).

Nakayama, K. and Loomis, J. M. 1974
Optical velocity patterns, velocity sensitive neurons and space perception
Perception 3, 63-80.

Nash, J. C. 1979
Compact Numerical Methods for Computers
New York: Wiley.

Prazdny, K. 1980
Egomotion and relative depth map from optical flow
Biological Cybernetics 36, 87-102.

Prazdny, K. 1981
Relative depth and local surface orientation from image motions
Technical Report TR-996, Computer Science Center, University of Maryland, College Park, MD 20742.

Regan, D., Beverley, K., and Cynader, M. 1979
The visual perception of motion in depth
Sci. Am. 241, 136-151.

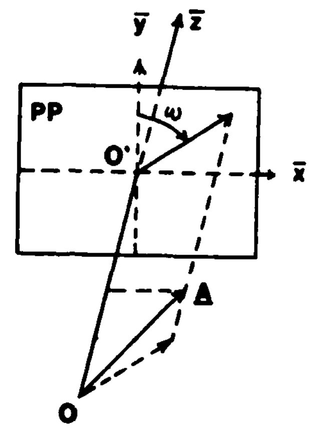

<u>Figure 1</u>.

Without loss of generality, the projection plane PP can be positioned at unit distance from the center of projection O, and parallel to the yz-plane. Any vector **A** can then be decomposed into a component parallel to the projection plane, and a component perpendicular to PP. O' is the center of the "retinal" coordinate frame (2D).

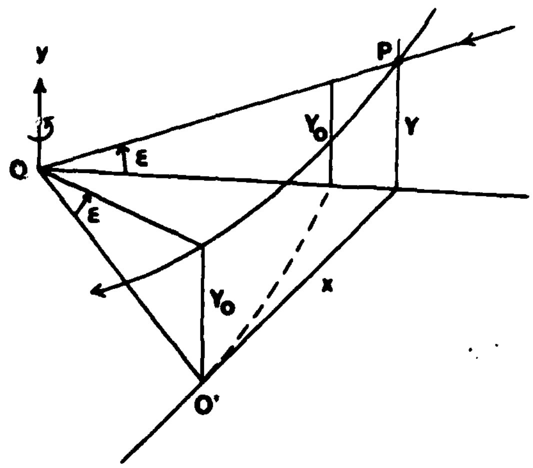

Figure 2.

When the observer rotates about the vector $\bar{y}$, the path described by an image element P on the image plane PP is a hyperbola.

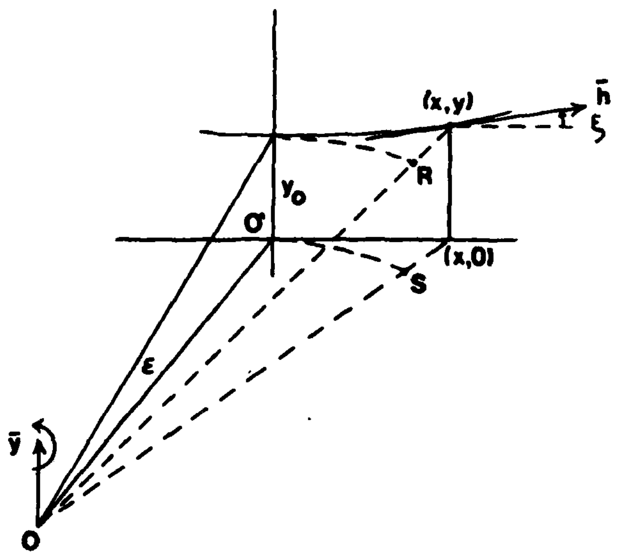

Figure 3.

The observer rotates about the line parallel to $\bar{y}$ through O. The direction of the velocity vector at P(x,y) is determined by $\tan\varepsilon$. The projections of the points R and S on PP have the same horizontal velocity components.

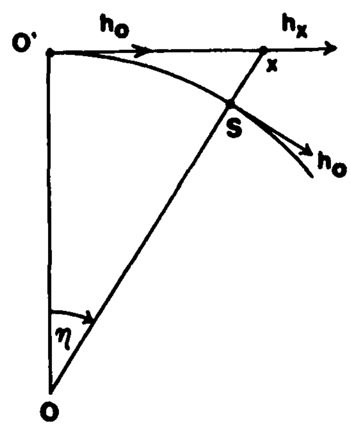

<u>Figure 4</u>.

The angular speed of a point x is $h_0 = \dfrac{h_x \sin(\pi/2 + \eta)}{\sqrt{x^2+1}}$, by definition ($|OO'| = 1$).
From this, h_x can be computed directly as a function of the parameter h_0.

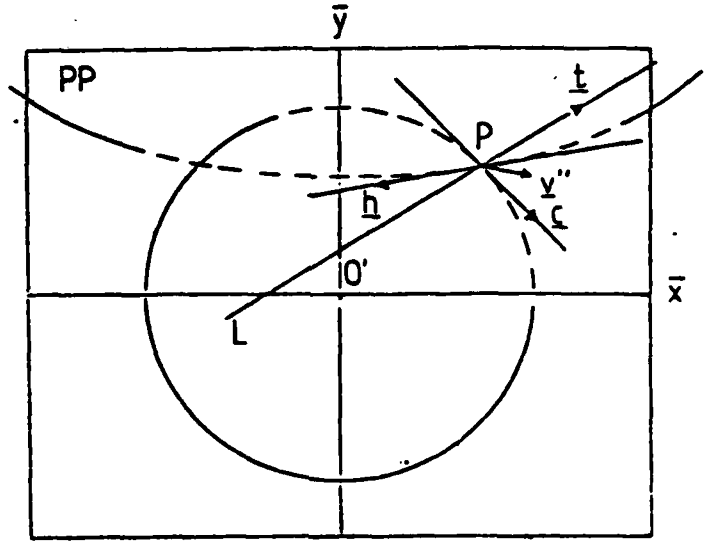

Figure 5.

An image velocity $\underline{v}''$ (on the planar projection surface PP) of a point P can be resolved into three components. The hyperbolic component $\underline{h}$ is due to the rotation of the ray about an axis (through O) in the projection plane (the angular velocity is a linear combination of $\bar{x}$ and $\bar{y}$). The circular component $\underline{c}$ is due to the rotation of the ray about an axis (through O) parallel to $\bar{z}$. The translational component $\underline{t}$ is the remaining vector which constrains the decomposition of $\underline{v}''$; $\underline{t}$ is constrained to be such that $\forall Q_i \in PP$: (L_i intersect in one common point). In the illustration above, the direction angle of the hyperbolic field is zero, i.e., the observer rotates only about a line parallel to the y-axis.

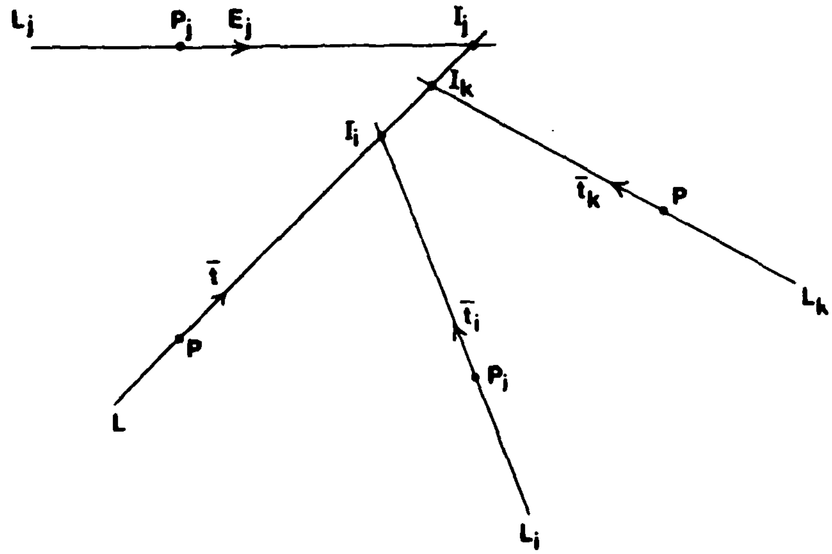

Figure 6.

To find the intersection of L with L_i (I_i), we have to solve for ℓ in

$$\underline{P}+\ell\underline{t}=\underline{P}_i+\ell_i\underline{t}_i$$

To obtain ℓ we multiply both sides by $\underline{t}_i'$, the perpendicular to $\underline{t}_i$. This yields

$$\ell=\frac{\underline{t}_i' \cdot (\underline{P}_i-\underline{P})}{\underline{t}\cdot\underline{t}_i}$$

where $\underline{P}$ and $\underline{P}_i$ are the (2D) position vectors on the projection plane. If $\underline{t}=(t_x,t_y)$ then $\underline{t}'=(-t_y,t_x)$.

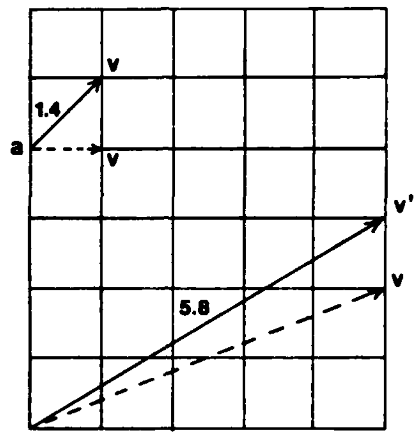

Figure 7.

The quantization error increases with decreasing vector magnitude. At b, an error of about 10 degrees of arc is made by representing v as v', while at a such a representation results in an error of 45 degrees of arc.

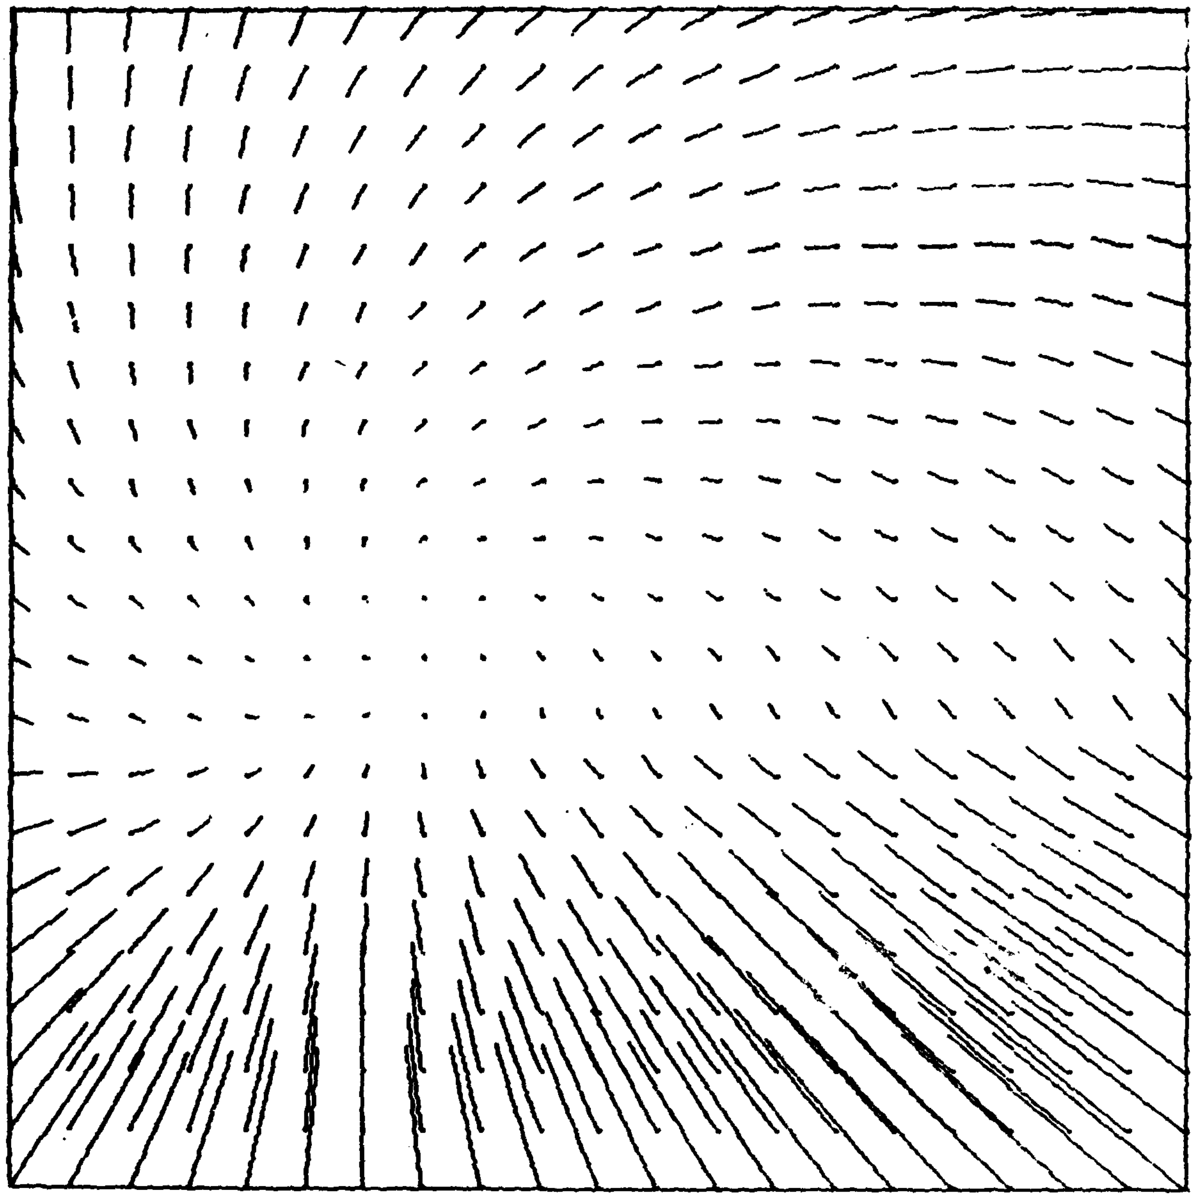

Figure 8.

(a) The positional velocity vector field generated by an observer "walking" backwards (on a horizontal ground plane) from a surface (plane) slanted 30 degrees (towards the observer) from the vertical.

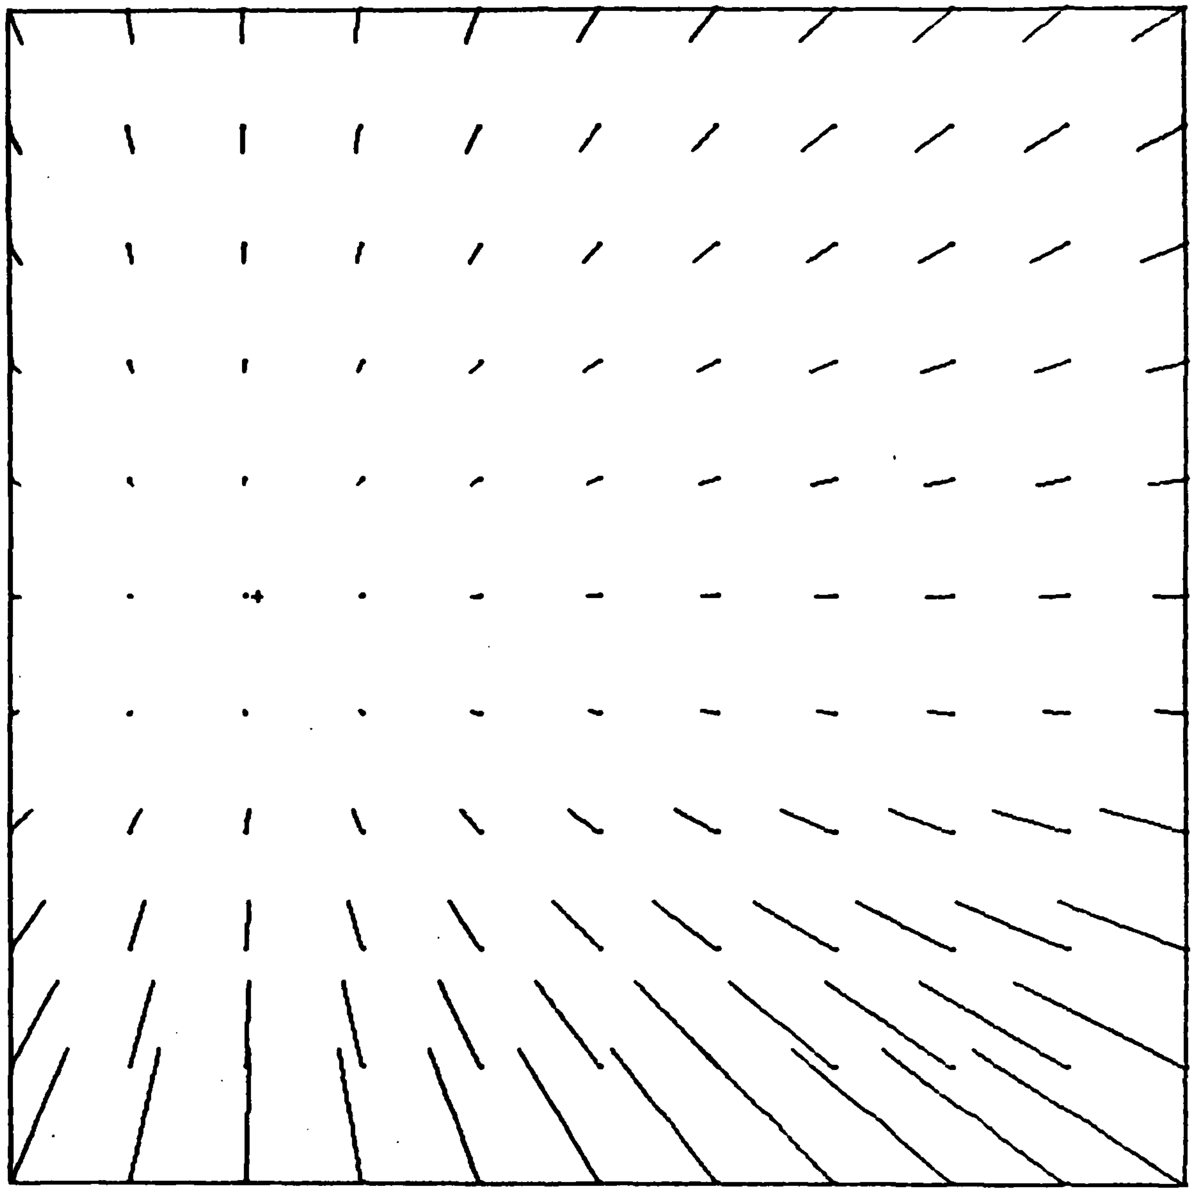

<u>Figure 8</u>.

(b) Corresponding translational field.
"+" denotes the FOE computed by the method (error=0).
The information about the (local) surface orientation and relative depth is contained <u>only</u> in the magnitudes of the velocity vectors; the direction of the vectors (and the FOE) depend only on the parameters of the relative motion.

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER	2. GOVT ACCESSION NO. AD-A109 566	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) Determining the instantaneous direction of motion from optical flow generated by a curvilinearly moving observer		5. TYPE OF REPORT & PERIOD COVERED Technical
		6. PERFORMING ORG. REPORT NUMBER TR-1009
7. AUTHOR(s) K. Prazdny		8. CONTRACT OR GRANT NUMBER(s) DAAG-53-76C-0138 DARPA Order-3206
9. PERFORMING ORGANIZATION NAME AND ADDRESS Computer Vision Laboratory, Computer Science Center, University of Maryland, College Park, MD 20742		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS
11. CONTROLLING OFFICE NAME AND ADDRESS U.S. Army Night Vision Laboratory Fort Belvoir, VA 22060		12. REPORT DATE February 1981
		13. NUMBER OF PAGES 27 (12) 30
14. MONITORING AGENCY NAME & ADDRESS(if different from Controlling Office)		15. SECURITY CLASS. (of this report) Unclassified
		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE

16. DISTRIBUTION STATEMENT (of this Report)

Approved for public release; distribution unlimited.

17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)

18. SUPPLEMENTARY NOTES

19. KEY WORDS (Continue on reverse side if necessary and identify by block number)

Image processing
Scene analysis
Optical flow
Time-varying imagery

20. ABSTRACT (Continue on reverse side if necessary and identify by block number)

A method is described capable of decomposing the optical flow into its rotational and translational components. The translational component is extracted implicitly by locating the focus of expansion associated with the translational component of the relative motion. The method is simple, relying on minimizing an (error) function of 3 parameters. As such, it can also be applied, without modification, in the case of noisy input information. Unlike

DD FORM 1 JAN 73 1473 EDITION OF 1 NOV 65 IS OBSOLETE

UNCLASSIFIED 403018 slt

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered) the previous attempts at interpreting optical flow to obtain information about the three-dimensional disposition of texture elements, the method uses only relationships between quantities on the projection plane. No 3D geometry is involved. Also outlined is a possible use of the method for the extraction of that part of the optical flow containing information about relative depth directly from the image intensity values, without extracting the "retinal" velocity vectors.